Comparing Internet Measurement Methods

Measuring internet performance continues to be meaningful for both consumers and service providers. A number of different measurement methodologies are typically used; however, the different methodologies frequently yield different results. It can be frustrating and confusing for consumers and service providers when results differ, even more so when results do not match their individual experience. The guide below is intended as an introduction to the common internet measurement methodologies, the strengths of each, their limitations, and some additional background on how Ookla chooses to measure internet performance.

Speedtest by Ookla uses consumer-initiated testing

Consumer-initiated tests are intentionally started by a person. Internet measurements occur at the times and in the places that are most relevant to the person taking the test. Each time a test is initiated, a snapshot of what the internet looks like in that place and time is recorded. When aggregated together, these individual experiences represent the typical internet performance, by carrier or provider, for a given location. Consumer-initiated tests can usually be performed on any device a person uses to access the internet, including desktop computers, laptop computers, smart phones and tablets. Speedtest is available on all these platforms, as well as others. Our dedicated test, which has been developed and refined for more than a decade, measures more than just internet speed and performance—it also reliably captures packet loss, jitter, errors, and other network conditions.

Consumer-initiated testing gives people the flexibility to initiate a test at any time and in any location, whether it be in homes, shopping centers, workplaces, parks, airports, as passengers in cars, on trains, or on buses, while in one place or on the move. When a critical mass of tests are taken, a comprehensive view of internet performance is achieved, including information about locations, times, service providers, and devices. While consumer-initiated testing has some limitations that are detailed below, Ookla believes there is no better mechanism to assess real-world internet performance.

For consumer-initiated testing to properly represent actual experience, the server network available for testing must be large and geographically dispersed. Speedtest leverages a vast testing infrastructure with over 6,000 servers in more than 190 countries, ensuring users will always be able to test to a local server regardless of where they are located. The network distance between the testing server and person taking the test is critical to performance measurement. By having multiple servers in every country and major city, Speedtest ensures an accurate view of performance without requiring long, or even international transit, to perform a test.

Of equal importance, the test volume must also be large enough to provide a statistically significant view across all desired dimensions (locations, service providers, technology types, and devices). Over 9.5 million tests are taken using Speedtest every day, which provides a truly

statistically significant sample size. All of these factors help ensure that Speedtest results accurately reflect the real-life performance people experience on the devices and the providers they actually use, from big cities to remote rural areas. As a result of the unparalleled number of tests taken using Speedtest, Ookla has no need to extrapolate data to fill in gaps because its applications collect information from every imaginable location and every type of device at all hours of the day.

Critics of consumer-initiated testing argue that, since the tests are taken by real people, the test conditions and timing are not randomized. Without laboratory grade randomization of test results, these critics believe the resulting data is skewed to reflect isolated conditions and times. In fact, Speedtest is designed to be able to test internet service that people actually experience, and with the expansive size of Ookla's server network, to return reliable performance data on the internet experience of a person in their specific location.

Following this same line of criticism, it is often assumed that consumer-initiated testing is utilized when people are having problems with their internet, meaning tests capture performance at its worst and do not reflect the complete range of speeds a person receives. We find that people use Speedtest at a widely diverse number of times and in an equally diverse set of situations, including when their internet is performing well. Massive testing volume on Speedtest applications also helps average out Ookla's aggregated findings by ensuring that outliers do not disproportionately influence results. Furthermore, the technology behind Speedtest is purpose-built to test the full potential of an internet connection. This technology, coupled with all the local server options available when using Speedtest, means that all test results truly capture the highest available potential of a network connection at any given point in time.

Another common criticism of consumer-initiated testing is that individuals, or even companies, intentionally or not, can skew the overall results for a given service provider. Without proper systems in place, especially when total test volume is low, this skewing effect could happen. To account for this, Ookla creates samples from all individual tests taken with Speedtest. Each sample represents the cumulative test results for each unique device/user per location, per calendar day. This improves accuracy of our data and reduces bias from things like repeated testing. Incidences of tampering with test volume or conditions are extremely rare, and if detected, are dealt with swiftly on both the Speedtest applications and data collection processes, as well as through discussions with the parties involved, if appropriate.

While Ookla has taken thorough steps to completely diminish or greatly reduce the impact of any negative side-effects born from our consumer-initiated testing, it is important to note that not all testing services that utilize consumer-initiated testing have taken the same steps. Other services also lack the same technology, server network, or testing volume to overcome the limitations.

Other measurement methodologies

Drive testing

Developed during the early days of cellular telephones, drive testing initially measured call quality and signal performance for car phone users. Drive testing was later adapted to measure

cellular internet performance and is still used today by many mobile providers to gauge the strength of their cellular networks.

Proponents of drive testing often cite that the process simulates precise lab testing conditions and is therefore a good way to determine the theoretical capacity and limits of a given cellular network. These conditions make drive test data particularly interesting for companies looking to map their network and understand regional opportunities for investment.

A primary limitation for drive testing is that the method still relies on driving lab test vehicles around vast geographic areas. This means data is only collected in areas where cars can drive. For example, roads comprise just six-tenths of one percent of the total land area in the contiguous United States according to the American Road and Transportation Builders Association estimates. As such, drive testing only from roads excludes the vast majority of locations in the US where consumers access the internet. This limitation is shared in countries globally.

Even though drive testing only covers areas with roads, it still requires a lot of time and resources to fully test an area. To account for this, only a few tests are usually taken in a given location and specific roadways are often not re-tested on a frequent basis. For instance, most areas may only be tested once or twice a year. Even if a company possessed a fleet of vehicles numbering in the hundreds, they still face significant logistical challenges that make it unrealistic for them to produce a technically comprehensive data set that reflects current consumer reality.

Additionally, drive testing cannot measure how connectivity changes throughout the day. Test vehicles drive through an area at a specific time, not all areas at all times of the day. This means drive testing has to extrapolate a small set of results over an extremely limited time frame when representing performance in a given geographic area. That also means that drive testing results are heavily influenced by momentary, abnormal fluctuations in network performance that might just happen at the single moment a drive test occurs. Drive testing relies on a limited set of fully optimized test hardware that does not represent the tremendous variety of devices that actual consumers use.

There are also concerns about gaming of test results when it comes to drive testing. Companies frequently are either aware of, or work with drive testers to define the routes, timing and devices being used during the drive testing. Pre-existing knowledge of these factors can create opportunities for manipulation of network conditions, such as increased infrastructure investments and tuning along roadways where tests are being conducted.

Background testing

Background testing applications, once installed onto a device, run periodic background tests whenever the device is on. The tests run at any and all times.

One of the greatest arguments in favor of background testing is that the random nature of testing means that data is collected in places where people either do not use the internet or think to initiate a test, which potentially leads to more dispersion of data points. Thoroughly dispersed data points are indeed useful for building coverage maps. At the very least, background testing generates significant amounts of data because testing is continuously

occurring. It is also favored for providing long-term trend information about performance and other measurements, like signal strength.

In practice, background testing is typically utilized by apps or vendors who do not have a large enough user base to rely on consumer-initiated testing results for their analysis alone. Background testing allows these apps and vendors to collect more data from a smaller user base.

The random timing of tests also leads to the largest criticism of background testing—that the tests do not coincide with times when a person would actually access the internet, which makes the data collected irrelevant. For example, background tests can occur when devices are in pockets, in the bottom of backpacks, in gym lockers, inside a nightstand, or in a sub-level parking garage, each of which is a situation where data is affected by signal attenuation. Furthermore, few background tests are even capable of thoroughly measuring internet speed.

Passive testing

Passive testing is a form of background testing that monitors the network usage of other applications on the mobile device itself. It can also measure the network usage of the device overall. Such testing methods focus only on the throughput afforded to a specific service or application, which can be interesting to app and device manufacturers.

Passive testing's big benefit comes from its ability to determine the performance of the apps and services a person uses on their device. Passive test results are most indicative of the internet experience a person receives when using a specific app or service, such as a video streaming app. Given that passive testing at its very nature runs in the background and is not a dedicated, standalone test, any data collected through background testing does not and cannot measure the full throughput of an internet connection.

The functions required to perform passive testing are not permitted to run on Apple's iOS platform, so the resulting data is limited to Android devices. This sample size is further limited since passive testing relies on small, pre-selected consumer panels. Those interested in passive testing should be aware that these limitations yield results that are liable to not be fully representative of how their actual internet connection performs.

Side-by-side benefits

	Measures anywhere	Tests all devices	Reflects network experience (on consumer SIM cards)	Tests all times of day	Replicable test patterns	Significantly large user base
Consumer-initiated testing with Speedtest	✓	✓	✓	✓	✓	✓

	Measures anywhere	Tests all devices	Reflects network experience (on consumer SIM cards)	Tests all times of day	Replicable test patterns	Significantly large user base
Drive testing					✓	
Background testing	✓	✓	✓	✓	✓	
Passive testing	✓	✓		✓	✓	

The Ookla approach

While arguments can be made in defense of all forms of internet testing, Ookla firmly stands behind the advantages of consumer-initiated testing as the most comprehensive and accurate view of internet performance and accessibility. The choice to use consumer-initiated testing stems from the fact that, out of all forms of testing available, it is the closest reflection of the internet experience people have in their daily lives and the way networks are performing for their users. Since the internet is a utility above all else, we believe that data and analysis about internet performance should be rooted in the context of these collective experiences.

Ziff Davis

IGN
AskMen
PCMag
Offers.com
ExtremeTech
Geek
Toolbox

OOKLA

Privacy Policy
Terms of Use